Data Analytics at NERSC

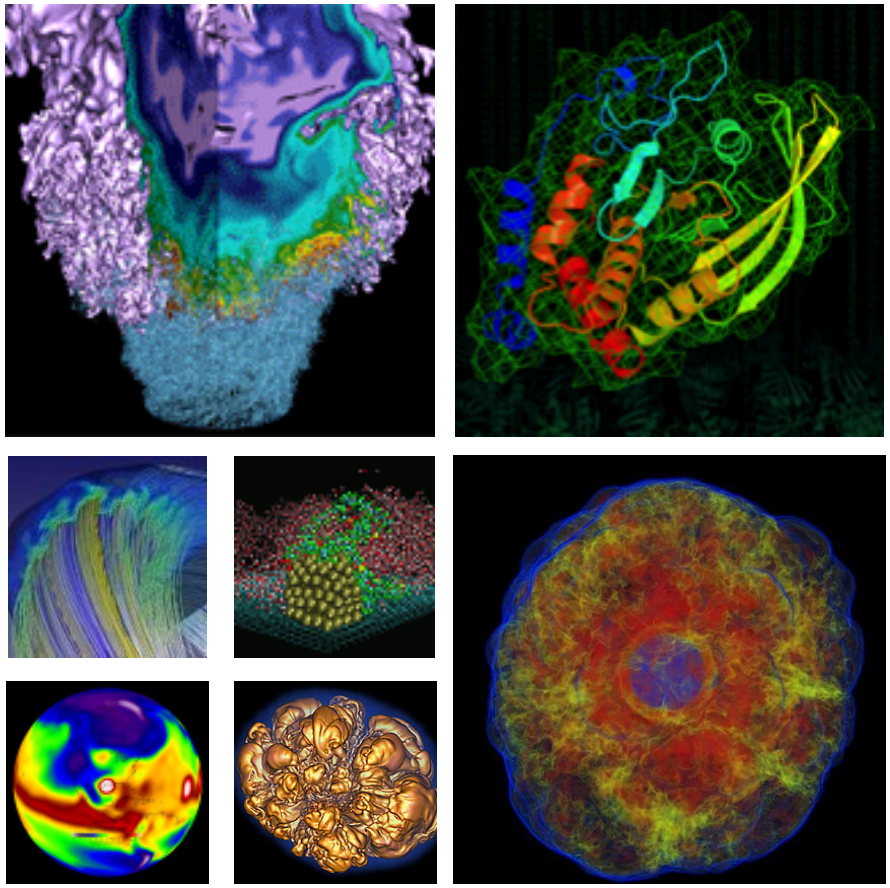

Joaquin Correa
JoaquinCorrea@lbl.gov
NERSC Data and Analytics Services

NERSC User Meeting
August, 2015

NERSC 40 YEARS at the FOREFRONT 1974-2014

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

Data analytics at NERSC

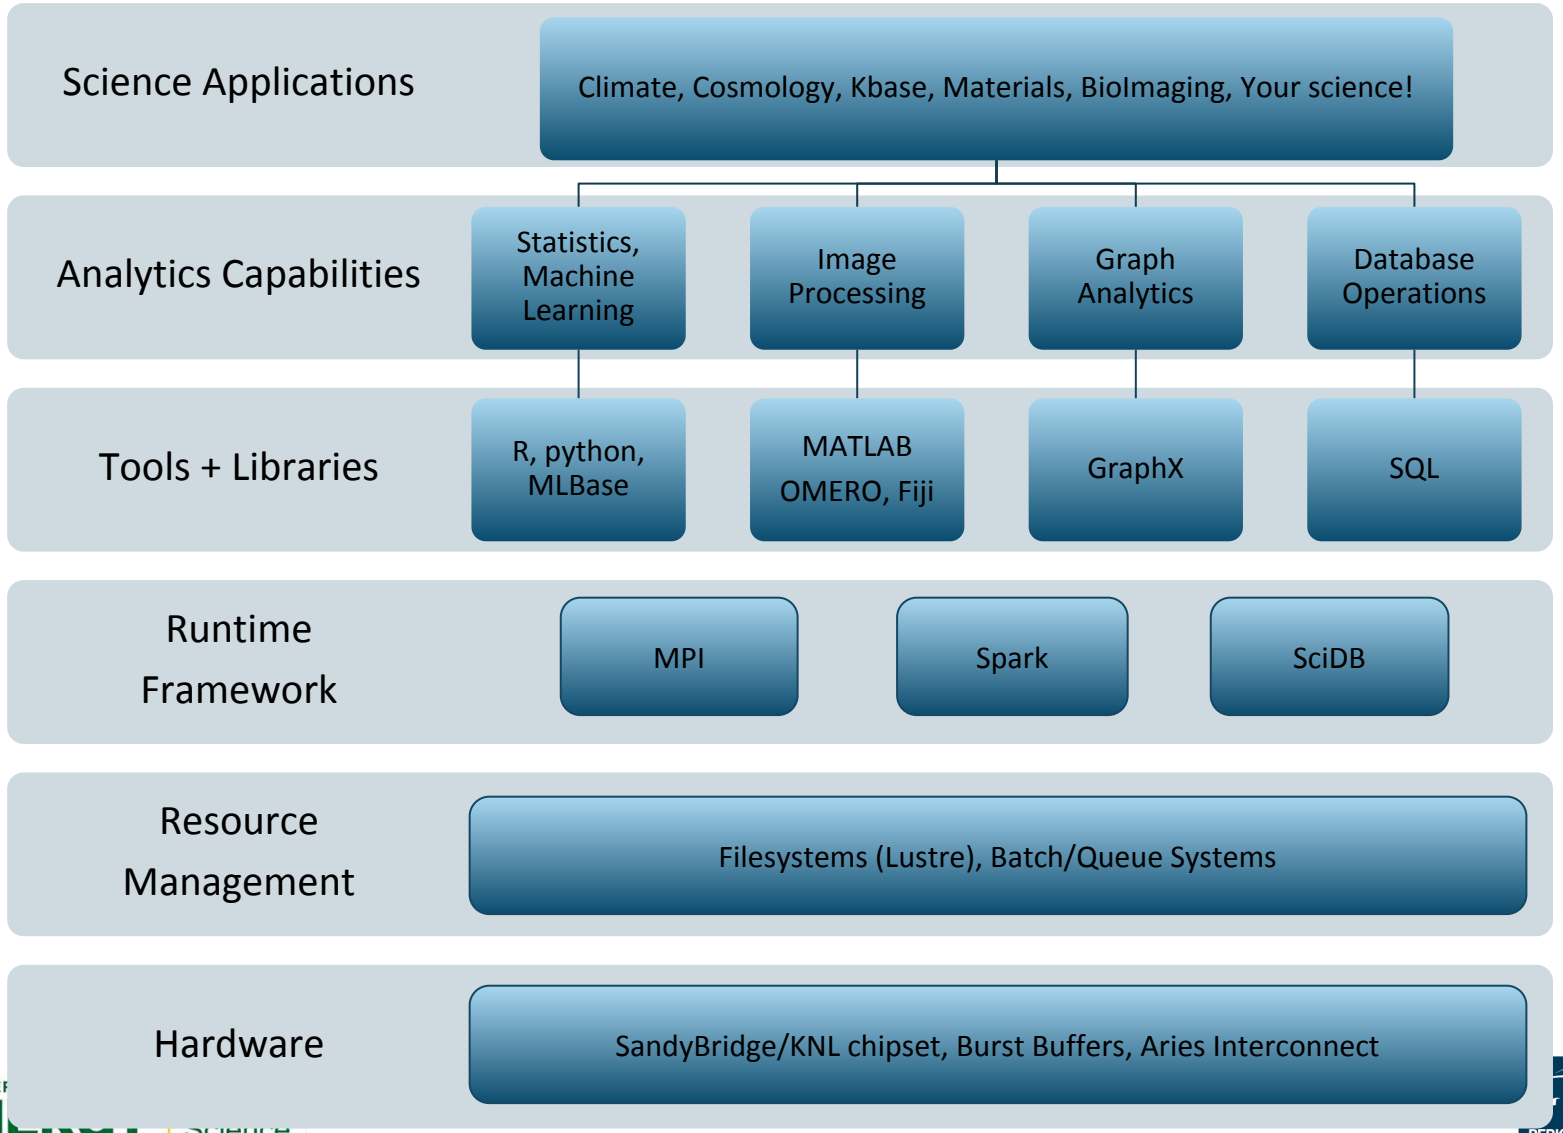

Data analytics at NERSC

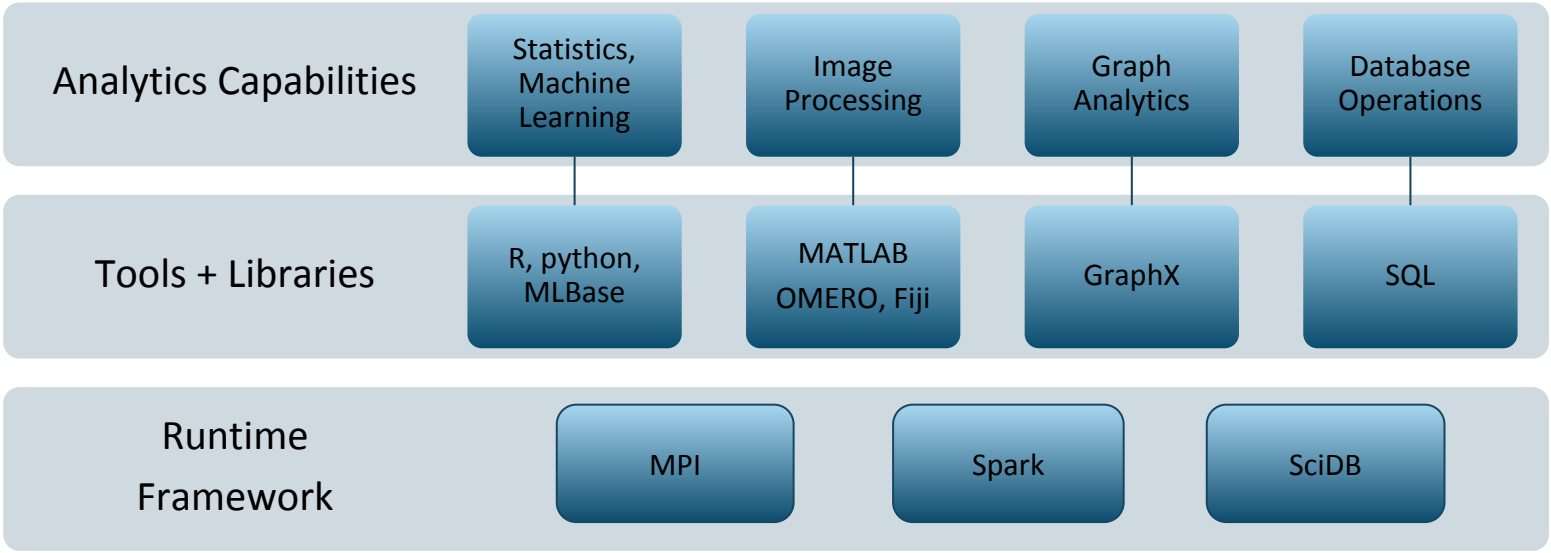

Talk Overview

- Data analytics tools
- Data insight
- Scale your analysis

Talk Overview

- **Data analytics tools**
- Data insight
- Scale your analysis

- R is a language and environment for statistical computing and graphics. It provides a wide variety of statistical tools, such as linear and nonlinear modelling, classical statistical tests, time-series analysis, classification, clustering, graphics, and it is highly extensible.

```
# Goal: To do `moving window volatility' of returns.

library(zoo)

# Some data to play with (Nifty on all fridays for calendar 2004) --
p <- structure(c(1946.05, 1971.9, 1900.65, 1847.55, 1809.75, 1833.65, 1913.6, 1852.65

# Shift to returns --
r <- 100*diff(log(p))
head(r)
summary(r)
sd(r)

# Compute the moving window vol --
vol <- sqrt(250) * rollapply(r, 20, sd, align = "right")

# A pretty plot --
plot(vol, type="l", ylim=c(0,max(vol,na.rm=TRUE)),
     lwd=2, col="purple", xlab="2004",
     ylab=paste("Annualised sigma, 20-week window"))
grid()
legend(x="bottomleft", col=c("purple", "darkgreen"),
       lwd=c(2,2), bty="n", cex=0.8,
       legend=c("Annualised 20-week vol (left scale)", "Nifty (right scale)"))
par(new=TRUE)
plot(p, type="l", lwd=2, col="darkgreen",
     xaxt="n", yaxt="n", xlab="", ylab="")
axis(4)
```

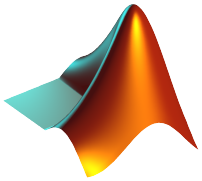

- MATLAB is a technical computing language that integrates computation, visualization, and programming in an easy-to-use environment where problems and solutions are expressed in familiar mathematical notation.

Toolboxes:

MATLAB	16
Image Processing	2
Neural networks	1
Optimization	2
Parallel computing	2
Signal processing	1
Statistics	2
Compiler	1

Mathematica

- Mathematica is a fully integrated environment for technical computing. It performs symbolic manipulation of equations, integrals, differential equations, and most other mathematical expressions. Numeric results can be evaluated as well.

- Fiji Is Just ImageJ - Fiji is an image processing package. It can be described as a "batteries-included" distribution of ImageJ (and ImageJ2), bundling Java, Java3D and a lot of plugins organized into a coherent structure.

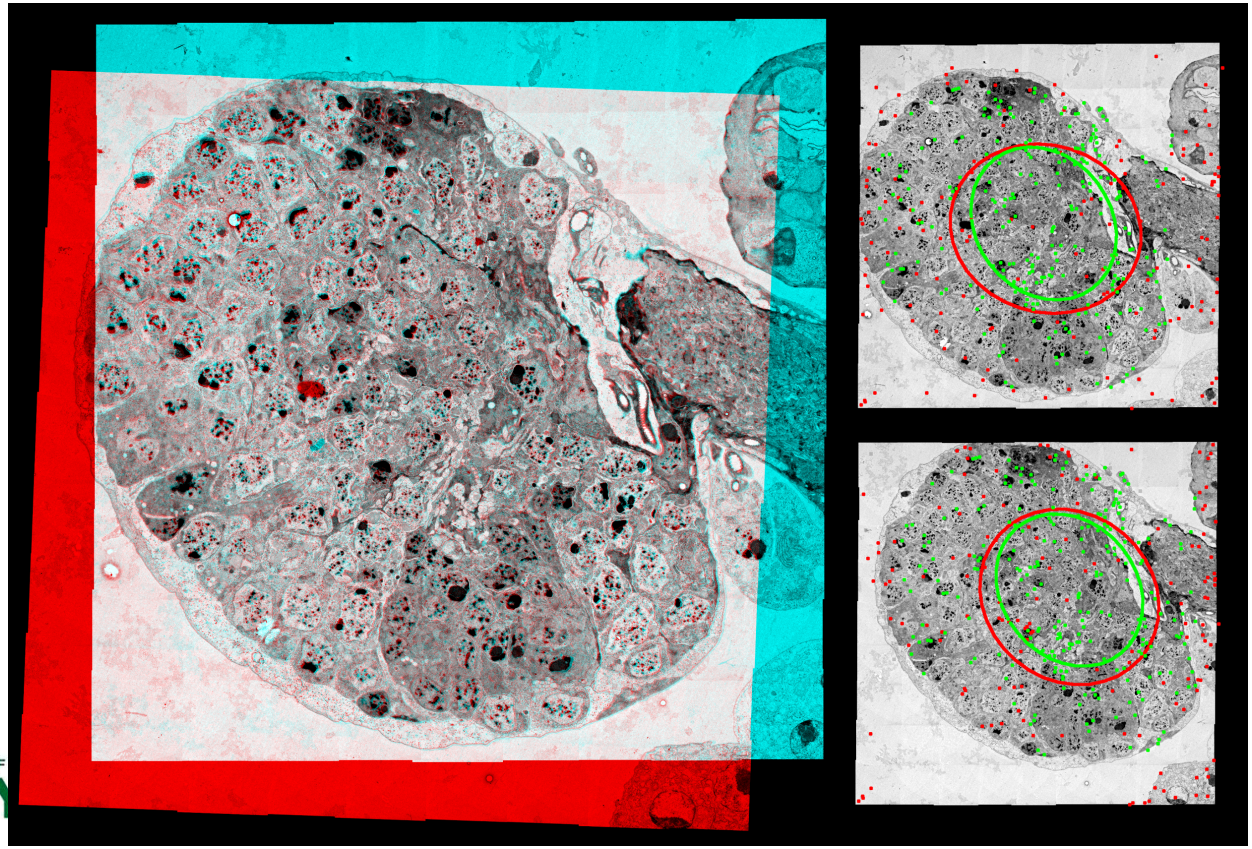

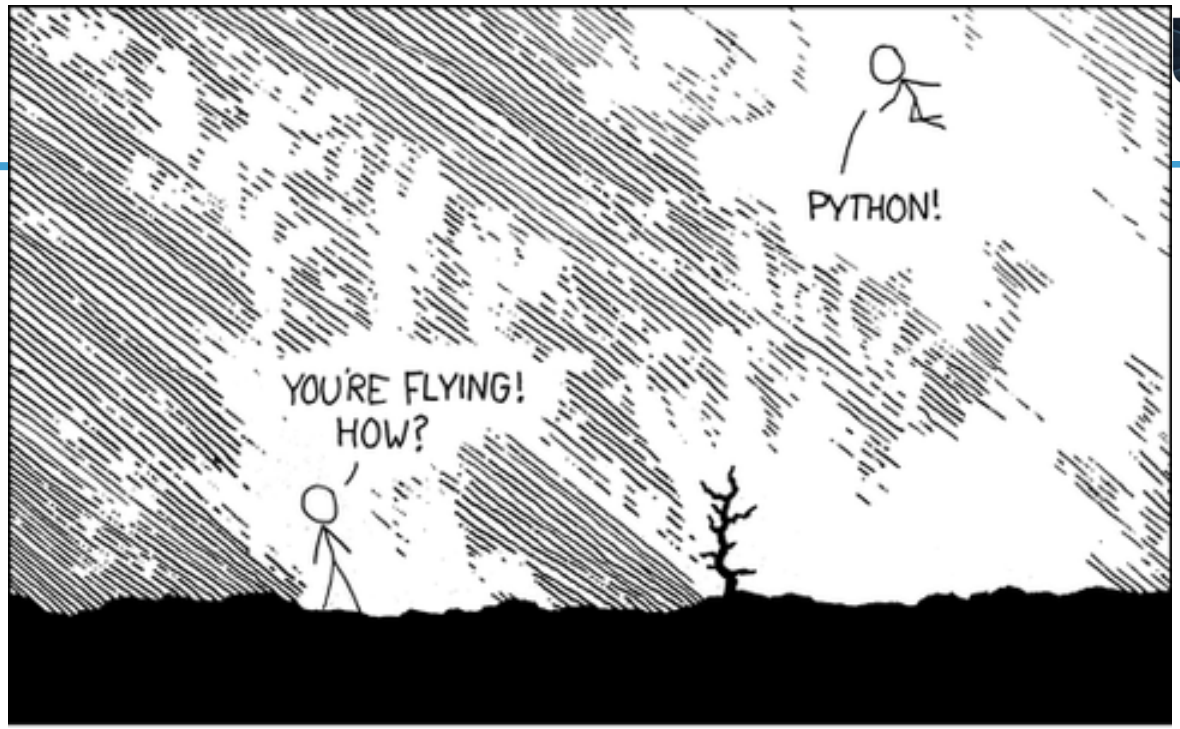
PYTHON!
YOU'RE FLYING! HOW?


I LEARNED IT LAST NIGHT! EVERYTHING IS SO SIMPLE!
HELLO WORLD IS JUST print "Hello, world!"


I DUNNO... DYNAMIC TYPING? WHITESPACE?
COME JOIN US! PROGRAMMING IS FUN AGAIN! IT'S A WHOLE NEW WORLD UP HERE!
BUT HOW ARE YOU FLYING?

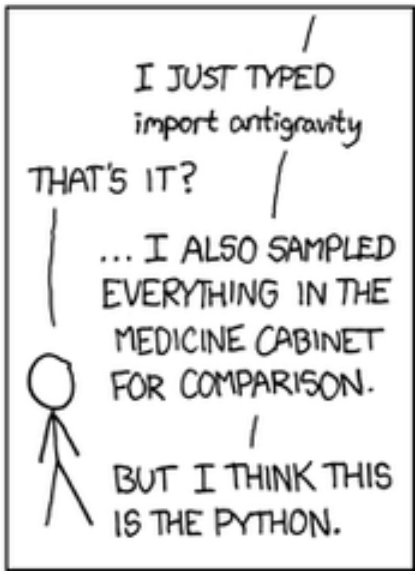
I JUST TYPED import antigravity
THAT'S IT?
... I ALSO SAMPLED EVERYTHING IN THE MEDICINE CABINET FOR COMPARISON.
BUT I THINK THIS IS THE PYTHON.

- Scientific Computing Tools for Python
  - NumPy
  - The SciPy library
  - Matplotlib
  - pandas
  - SymPy
  - IPython
  - nose
  - Cython
  - Scikits
  - h5py
  - mpi4py

Deep learning at NERSC

- neon is an easy to use, python-based scalable Deep Learning library.

Deep Learning has recently achieved state-of-the-art performance in a wide range of domains including images, speech, and text. It is seeing adoption in the HPC community as a tool for large-scale data processing.

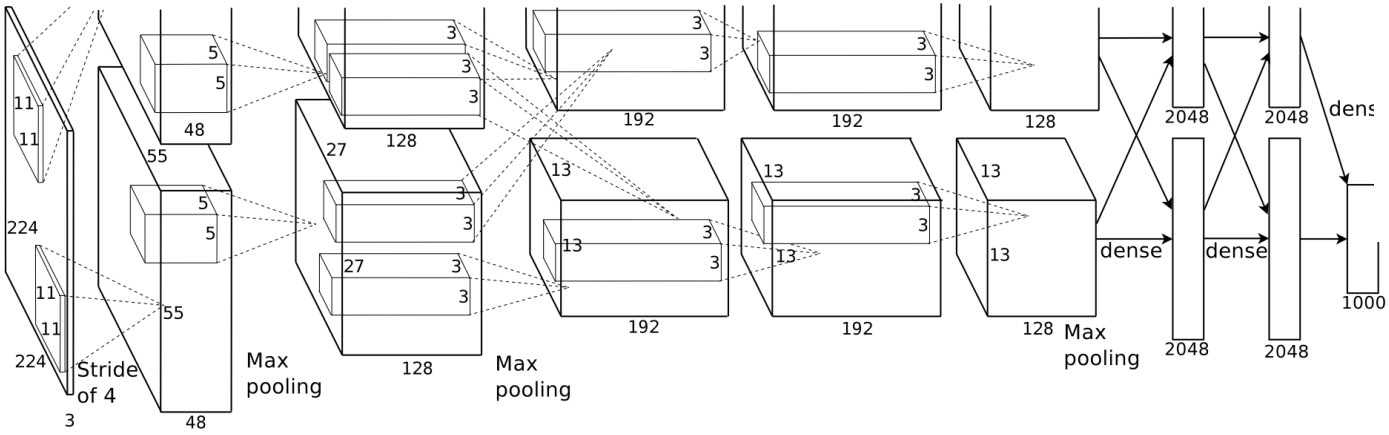

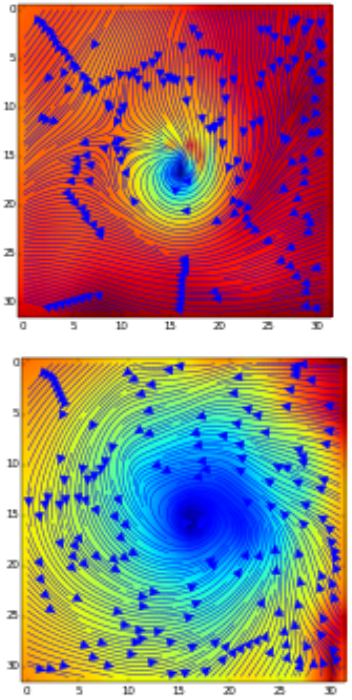

Talk Overview

- Data analytics tools
- **Data insight**
- Scale your analysis

R and RStudio Service(Beta)

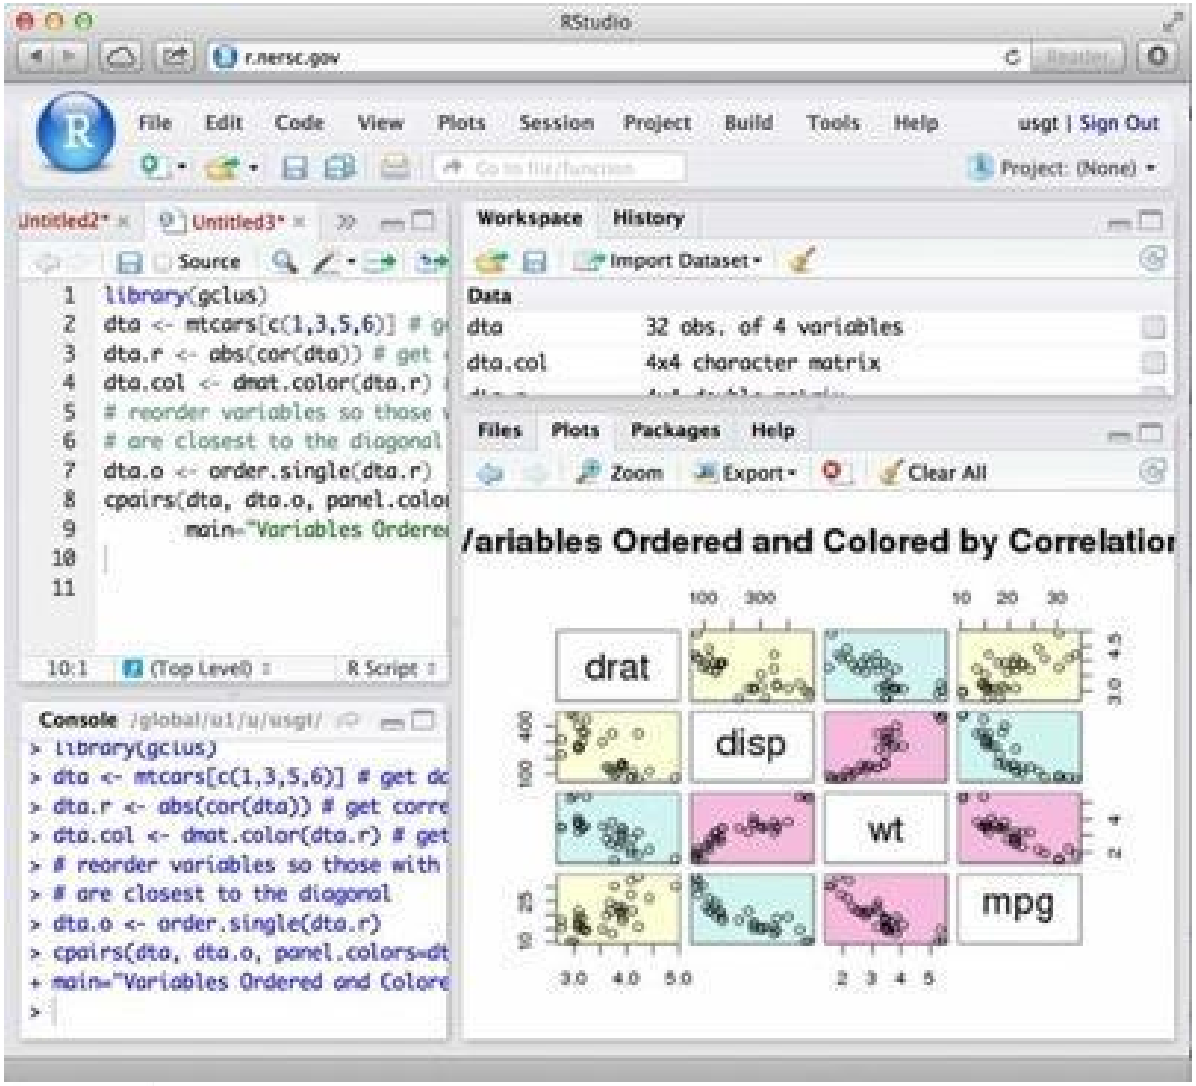

iPythonNotebook Service(Beta)

IP[y]: Notebook Lecture_3_Scipy (autosaved) Logout

File Edit View Insert Cell Kernel Help

Markdown Cell Toolbar: None Env: np18py27-1.9

In [6]:
```
x = linspace(0, 10, 100)

fig, ax = subplots()
for n in range(4):
    ax.plot(x, jn(n, x), label=r"$J_%d(x)$" % n)
ax.legend();
```

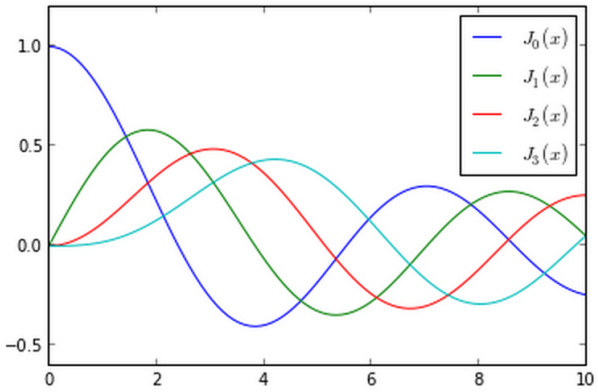

In [7]:
```
# zeros of Bessel functions
n = 0 # order
m = 4 # number of roots to compute
jn_zeros(n, m)
```

Out[7]:
```
array([  2.40482556,   5.52007811,   8.65372791,  11.79153444])
```

U.S. DEPARTMENT OF ENERGY | Office of Science

ipython.nersc.gov

BERKELEY LAB

Talk Overview

- Data analytics tools
- Data insight
- **Scale your analysis**

- Apache Spark is a fast and general-purpose cluster computing system. It provides high-level APIs in Java, Scala, Python and R, and an optimized engine that supports general execution graphs.

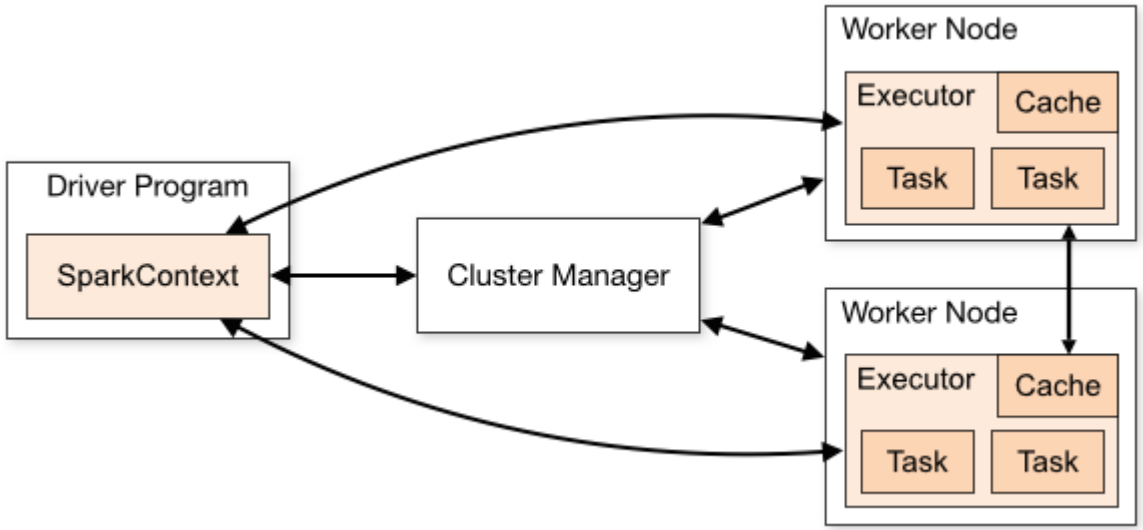

- Easy to Prototype
  - Interactive shells make it easy to explore your data
  - Interactively debug your analyses
  - Works with iPython notebooks

- Easy to Run
  - High-Level API using map-reduce paradigm
  - Implement all your analyses in a few lines of
  - Python
  - Scala
  - Java

Computation Type	Spark Implementation
Machine Learning	MLlib, Spark ML
Graph Computations	GraphX
Database Operations	Spark SQL
Streaming Analysis	Spark Streaming
Your Own Custom Analysis	Using Spark's Built In Functions

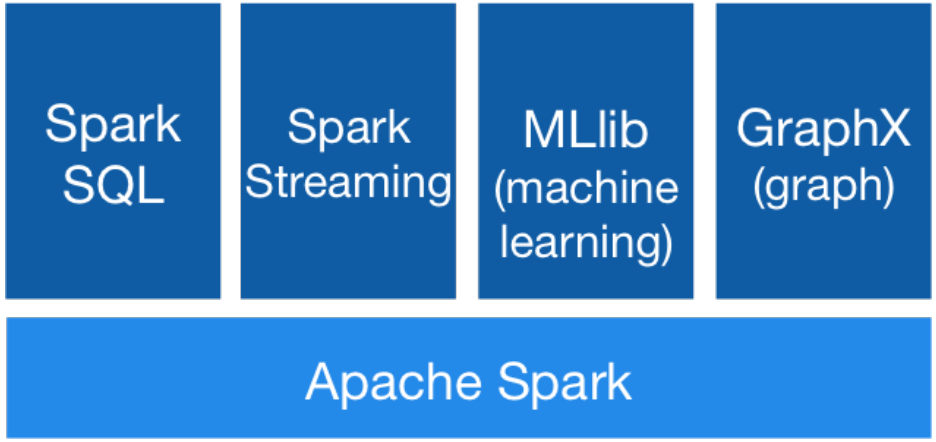

Computation types can be combined seamlessly all in the same piece of code!

SciDB For High Usability Big Data Analytic

- **Why?** It's painful to manage and analyze terabytes of data. Need a unified solution that's easy to use.
- **What?** SciDB is a parallel database for array-structured data, great for **Terabytes** of:
  - **Time series, spectrums, imaging**, etc
- The greatest benefit of SciDB is:
  - **Usability**: Use HPC hardware without learning parallel programming and parallel I/O.

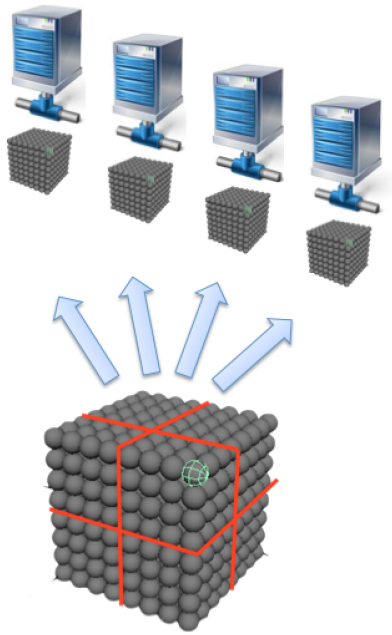

SciDB Distribute a big array on many nodes

NERSC Data Analytic Services

Big and Diverse Computing Facility
6000+ Users, 700+ Projects
3+ PetaFlops (20+pf more coming)
50+ PB Storage

Spark
PostgreSQL
SciDB
MySQL
SQL
R
IP[y]:
NOMACHINE
VOLTDB

Production Data Services

Science Engagement

Thank you.

Cori: Unified architecture for HPC and Big Data

- **64 Cabinets of Cray XC System**
  - 50 cabinets 'Knights Landing' ***manycore*** compute nodes
  - 10 cabinets 'Haswell' compute nodes for ***data partition***
  - ~4 cabinets of Burst Buffer
  - 14 external login nodes
  - Aries Interconnect (same as on Edison)
- **Lustre File system**
  - 28 PB capacity, 432 GB/sec peak performance
- **NVRAM "Burst Buffer" for I/O acceleration**
- **Significant Intel and Cray application transition support**
- **Delivery in mid-2016; installation in new LBNL CRT**

NERSC Systems

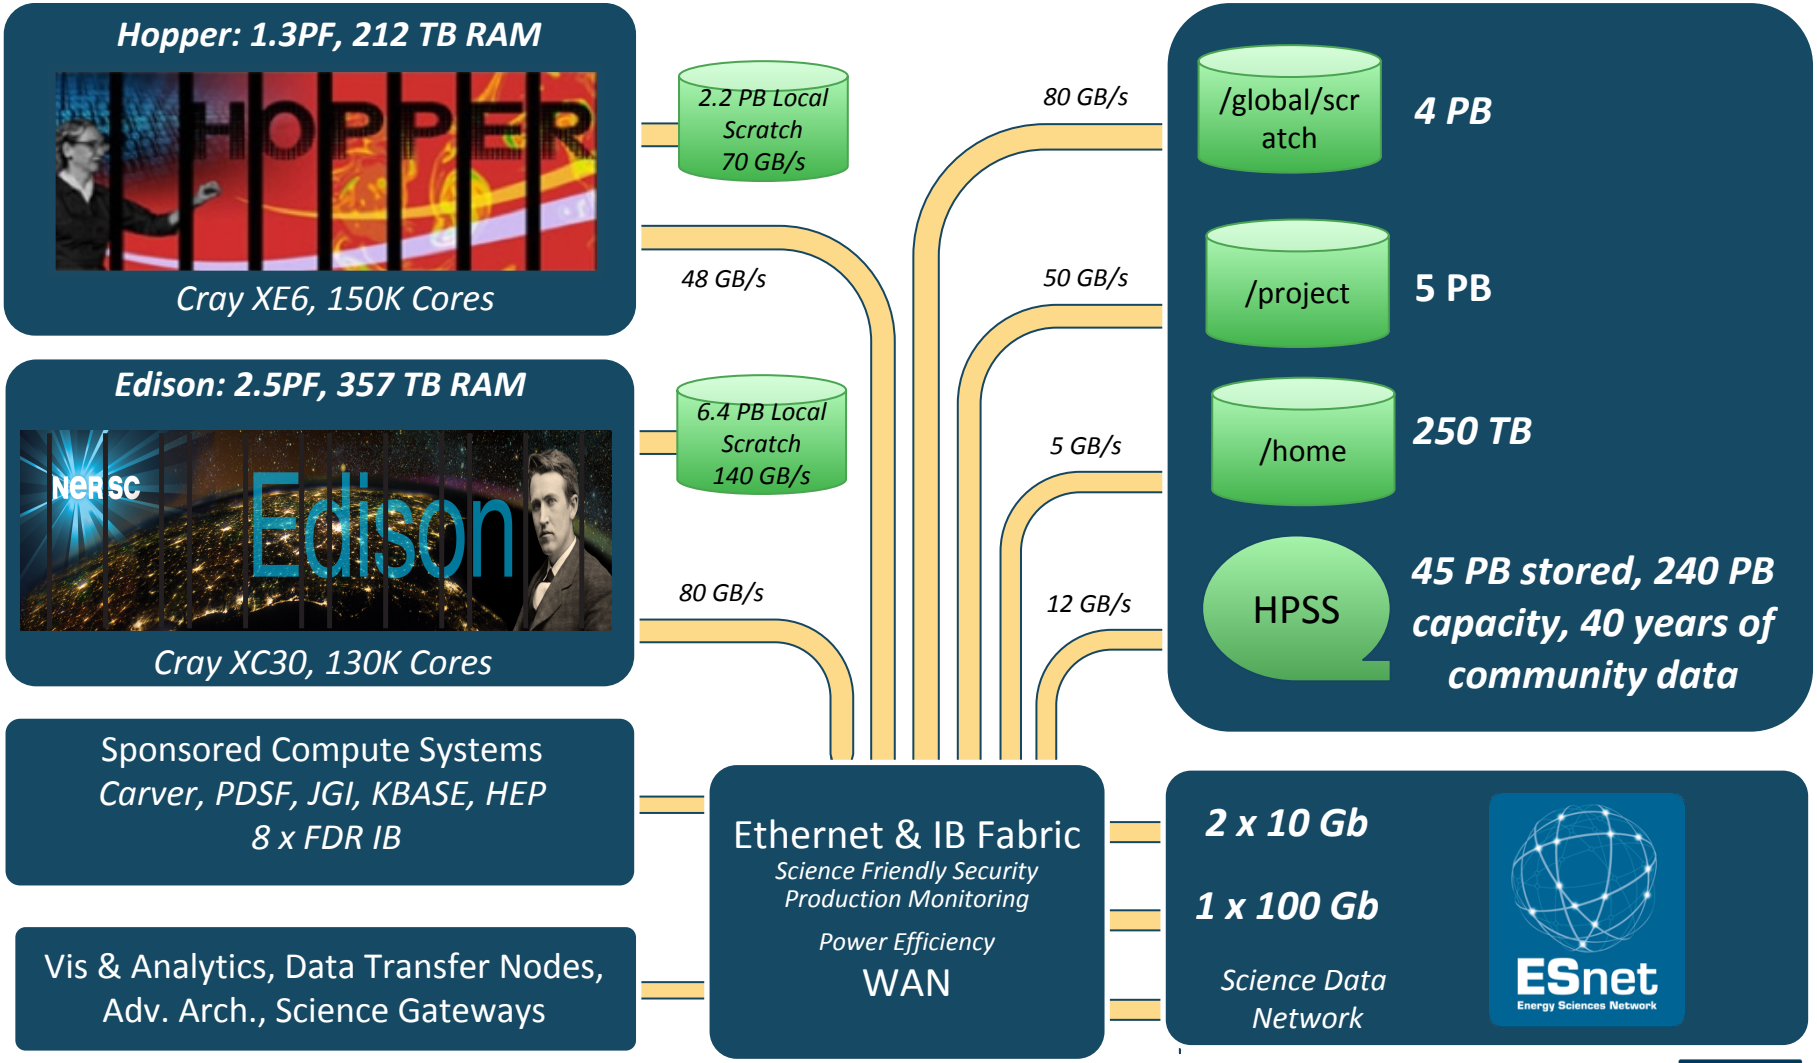

5 V's of Scientific Big Data

Science Domain	Variety	Volume	Velocity	Veracity
Astronomy	Multiple Telescopes, multi-band/spectra	O(100) TB	100 GB/night – 10 TB/night	Noisy, acquisition artefacts
Light Sources	Multiple imaging modalities	O(100) GB	1 Gb/s-1 Tb/s	Noisy, sample preparation/acquisition artefacts
Genomics	Sequencers, Mass-spec, proteomics	O(1-10) TB	TB/day	Missing data, errors
HEP: LHC, Daya Bay	Multiple detectors	O(100) TB – O(10) PB	1-10 PB/s reduced to GB/s	Noisy, artefacts, spatio-temporal
Climate	Simulations Multi-variate, spatio-temporal	O(10) TB	30-70 GB/s	'Clean', need to account for multiple sources of uncertainty

DOE Facilities are Facing a Data Deluge

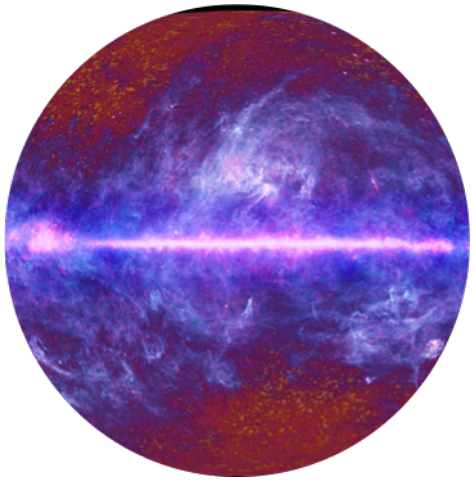

Astronomy

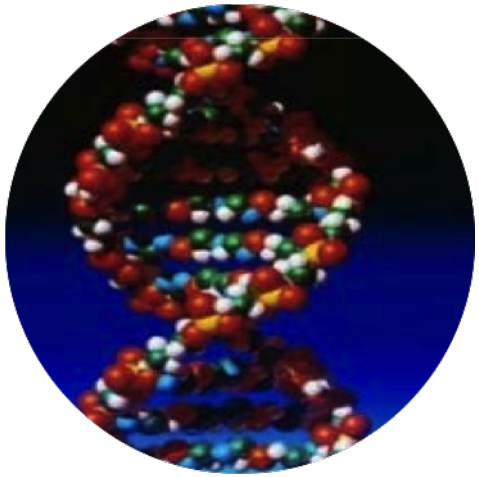

Genomics

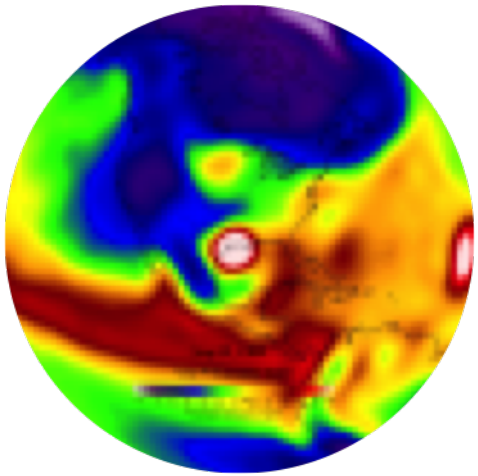

Climate

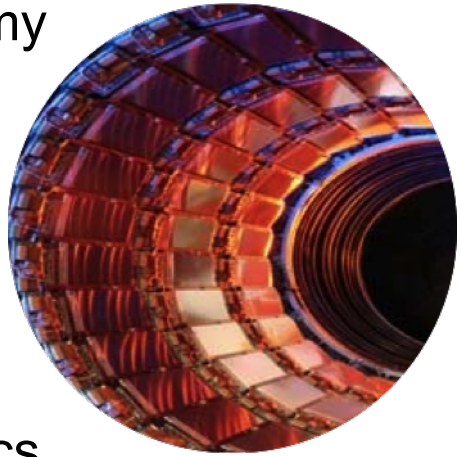

Physics

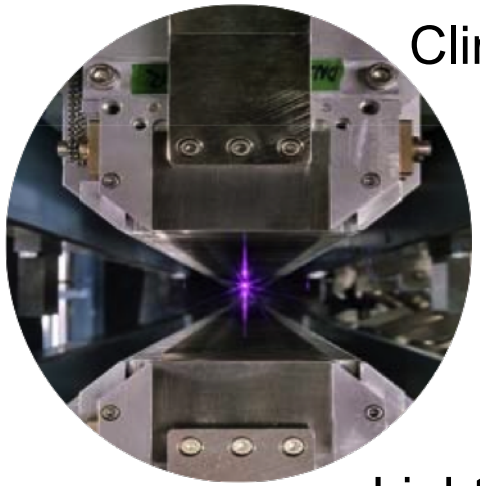

Light Sources